

June 12, 2023

Response to: NTIA AI Accountability Policy Request for Comment

Artificial intelligence technologies are quickly being integrated into a wide range of consumer technology with insufficient attention paid to the potential risk and impact of these systems. The vast majority of these systems are being built on top of algorithms and data collected and managed by the world's largest technology companies, several of which have laid off their ethics teams¹ or whose executives have campaigned against researchers criticizing their work².

AI accountability requires stricter guidelines on the capabilities and suggested use of the technology, as well as protections for data collection and model training. Companies building or selling software that utilizes artificial intelligence technologies should practice transparency, accountability, and ethical participation in disclosing training data and language model sources, and in the capabilities of their products and services.

Transparency - Companies who are publishing APIs for use by software developers to build other products and services should be required to provide a report to their customers that includes a detailed description of the sources and types of data being used in training their models. If the product or service is built using a publicly available, open source model, this should be disclosed. Specific regulatory guidance that is written for these companies should take into consideration the size and revenue of companies to ensure that there are ways to comply with transparency requirements that do not prohibit the entrance of new competitors to the market in providing AI models and services.

Accountability

Companies who are utilizing generative AI/machine learning APIs in commercial software applications should have clearly defined rules in the event of a realized harm based on the use of the technology. These harms should be evaluated and decided upon by a cross-functional and multi-disciplinary group of technology builders, social scientists, academics, and members of marginalized communities most likely to be impacted by the downstream effects of the adoption and use of generative AI.

Companies building large language models available for use in commercial applications that meet any of the following criteria should be required to allow a third-party to audit the sources of their data, storage, and use. Specific regulatory guidance should be written with scaling requirements that become more intensive relative to the size of the company (by revenue) or use.

Sample criteria³ might include:

- Language models that have a significant number of parameters (for example, more than 13B)
- Language models that are partially or fully trained/fine-tuned on publicly available online data (examples of which may include, but are not limited to, source code made available under an open

¹ As AI booms, tech firms are laying off their ethicists, The Washington Post - March 20, 2023, <https://www.washingtonpost.com/technology/2023/03/30/tech-companies-cut-ai-ethics/>

² Bigshot chief scientist of major corporation can't handle criticism of the work he hypes, Columbia University Statistics - November 23, 2022, <https://statmodeling.stat.columbia.edu/2022/11/23/bigshot-chief-scientist-of-major-corporation-cant-handle-criticism-of-the-work-he-hypes/>

³ A note on regulating based on specific criteria for models - techniques to distribute and train smaller models are rapidly being developed. File size may not be a significant factor in future models, so any regulatory guidance that uses the size of a model as a proxy for regulatory evaluation should remain flexible as to account for the changing technological landscape, perhaps updated on a regular cadence (e.g. quarterly)

source license, Creative Commons licensed writing, public social media posts made by individual accounts, and public domain literature)

- Text-based Language models that are more than 5GB in size when quantized
- Models that meet a computational cost threshold (as measured in C/G/TPU cycles or USD) to train or fine-tune
- OSHA-like rules to protect workers who are training models and labeling data from the impact to mental health that arises from repeated exposure to harmful content

Additionally, companies providing APIs that interface with large language models should indicate that they are utilizing such models to their customers.

Ethical Participation

One significant challenge present with AI models is the difficulty in understanding how a model comes to a given output. There are many steps that result in machine-generated output that change based on the data used to train a model, fine-tuning steps that embed new data on top of an existing model, the prompt(s) used to dictate behavior of the output, and other restrictions or capabilities that developers might put into place to modify or attempt to control the use of the system.

Companies building AI services or products that utilize AI for sale to end customers should be required to articulate their approach to their use of the systems. Ethical participation can be measured through a variety of frameworks, but should include elements of an analysis of the organizational effort to provide transparency and hold themselves accountable for their systems, realized harms of their systems when misused or used for harm, and willingness to comply with requests for information about how these systems are built and created.

Closing thoughts

Within the scope of the NTIA AI Accountability request for comment, three recommendations are made above. However, limiting the accountability conversation to measurable impact of systems is a nuanced and complex situation that should further take into consideration the immense amount of centralized power that technology oligopolies hold in this space, the use of data without permission in training sets⁴, and human rights violations⁵ that are occurring as part of the manner in which these oligopolies are choosing to advance their technological authority.

We the undersigned ask the NTIA to deeply consider the larger issues of technological antitrust, individual copyright agency, and human rights protections as part of their accountability requirements for companies, while recognizing that the existing system adversely favors large corporations and regulation can inhibit further competition by smaller agencies, open source communities, and individuals. Through the lessons learned from a lack of regulation on the use of data collection and social media, and the more recent scrutiny on technology monopolization of our social graphs, the NTIA can make meaningful

Signed,

Olivia Erickson

Co-Founder, iamagoodbing.ai

Executive MBA Candidate, Columbia Business School

⁴ Artists are Suing Artificial Intelligence Companies, and the Lawsuit could Upend Legal Precedents Around Art
<https://www.artnews.com/art-in-america/features/midjourney-ai-art-image-generators-lawsuit-1234665579/>

⁵ Open AI and Sama Hired Underpaid Workers in Kenya to Filter Toxic Content for ChatGPT
<https://www.business-humanrights.org/en/latest-news/openai-and-sama-hired-underpaid-workers-in-kenia-to-filter-toxic-content-for-chatgpt/>

Zachary Fox

Co-Founder, iamagoodbing.ai

M Eifler

Founder, BlinkPopShift